arXiv:1607.05709v1 [math.ST] 19 Jul 2016

Multi-category Angle-based Classifier Refit

Guo Xian Yau[1] and Chong Zhang [2]

[1]Department of Mathematics
University of Waterloo
[2]Amazon.com, Inc.

Abstract

Classification is an important statistical learning tool. In real application, besides high prediction accuracy, it is often desirable to estimate class conditional probabilities for new observations. For traditional problems where the number of observations is large, there exist many well developed approaches. Recently, high dimensional low sample size problems are becoming increasingly popular. Margin-based classifiers, such as logistic regression, are well established methods in the literature. On the other hand, in terms of probability estimation, it is known that for binary classifiers, the commonly used methods tend to under-estimate the norm of the classification function. This can lead to biased probability estimation. Remedy approaches have been proposed in the literature. However, for the simultaneous multicategory classification framework, much less work has been done. We fill the gap in this paper. In particular, we give theoretical insights on why heavy regularization terms are often needed in high dimensional applications, and how this can lead to bias in probability estimation. To overcome this difficulty, we propose a new refit strategy for multicategory angle-based classifiers. Our new method only adds a small computation cost to the problem, and is able to attain prediction accuracy that is as good as the regular margin-based classifiers. On the other hand, the improvement of probability estimation can be very significant. Numerical results suggest that the new refit approach is highly competitive.

1 Introduction

Classification is one of the founding pillars in the statistical learning literature. It has achieved many successes on real applications in various scientific disciplines, such as artificial intelligence, cancer research, and econometrics. Given a training data set with both predictors and labels observed, one can obtain a classifier for future prediction. Typical examples of classification methods include logistic regression, Fisher's linear discriminant analysis, nearest neighbours, and many more. See Hastie *et al.* (2009) for a comprehensive review of existing methods.

For classification problems, the primary goal is to train the classifier with a high prediction accuracy. A classifier is optimal if its corresponding prediction error rate achieves the Bayes error rate, i.e., the smallest error rate among all possible classifiers. Besides prediction accuracy, class probability estimation is also very important in practice, especially for non-separable classification problems. In particular, when the density of two or more classes overlaps, there are random errors involved in class label prediction. In this case, class probability estimation can provide practitioners with valuable information on the confidence of classifying one subject into a certain category. For example, given the prediction that the stock market index would rise (as opposed to fall), the future investment strategies can be quite different if the estimated probabilities are 0.51 and 0.99 respectively.

In practice, if the dimension of the classification problem is low, traditional classifiers such as nearest neighbours are known to work well. Furthermore, class probability estimation can also be very accurate. However, recent technology development has enabled scientists to collect big data sets with very high dimensions. For example, microarray data can measure the expression levels of millions of genes simultaneously, whereas the number of observations is often limited. In this case, traditional learning methods are known to suffer from the potential danger of overfitting. That is, even if the classifier has a high prediction accuracy on the training data set, its generalization error on future testing data sets can be very large. To overcome this difficulty, many new classifiers have been proposed in the literature. For example, Cai and Liu (2011) proposed a generalization of Fisher's LDA method to handle high dimensional problems, and Lin *et al.* (2000) suggested the penalized logistic regression.

In this paper, we focus on margin-based classification methods. This family was originally introduced in the machine learning literature, and has become increasingly popular partly because of their ability to handle high dimensional classification problems. A typical margin-based classifier

solves an optimization problem in the *loss + regularization* form. Here the loss term is used to measure the goodness of fit of the classifier on the training data set. The regularization term prevents overfitting by imposing penalty on complexity of the functional space, aiming to obtain a classifier with accurate prediction on future testing data sets. Many existing margin-based classifiers have been proposed in this framework, including penalized logistic regression (Lin *et al.*, 2000), Support Vector Machines (SVM, Boser *et al.*, 1992; Cortes and Vapnik, 1995), import vector machines (Zhu and Hastie, 2005), Adaboost in Boosting (Freund and Schapire, 1997; Friedman *et al.*, 2000), Ψ-learning (Shen *et al.*, 2003), and large margin unified machines (Liu *et al.*, 2011).

Binary margin-based classifiers have been well studied. To address a multicategory classification problem, one can first consider a sequence of binary margin-based classifiers, such as the one-versus-one and one-versus-rest approaches. Although these methods are simple in concept and implementation, they have certain drawbacks. For example, the one-versus-one method has a potential tie-in-vote problem, and the one-versus-rest method can be inconsistent (Liu, 2007). Therefore, in this paper, we focus on simultaneous margin-based classifiers. In the literature, most simultaneous margin-based classifiers use k classification functions for a k-class problem. These methods often employ a sum-to-zero constraint in the corresponding optimization to reduce the parameter space and to ensure some theoretical properties such as Fisher consistency (also known as classification calibration, see Bartlett *et al.*, 2006). Many proposed classification methods follow this procedure. See, for example, Lee *et al.* (2004), Zhu and Hastie (2005), Wang and Shen (2007), Park *et al.* (2010), Liu and Yuan (2011), Zhang and Liu (2013), among others. However, this approach of using k functions and reducing to $k-1$ by the sum-to-zero constraint, despite its simplicity, can be inefficient. To circumvent this difficulty, Zhang and Liu (2014) proposed a new framework for simultaneous margin-based classification, namely, the angle-based classification, which only uses $k-1$ functions for classification. More details of angle-based classifiers are given in Section 2. Zhang and Liu (2014) showed that their method is free of the sum-to-zero constraint, hence for high dimensional problems, the angle-based method can enjoy faster computational speed and higher prediction accuracy. In this paper, we employ the angle-based framework for multicategory classification.

For margin-based classifiers, to estimate class probabilities, one commonly used approach is to explore their relationship with respect to the theoretical minimizer, and use the obtained classification functions for probability estimation (more details are given in Section 2). Many existing classifiers in the literature use this method to estimate class probabilities. See, for example, Lin

et al. (2000), Park *et al.* (2010), Liu *et al.* (2011), and Zhang and Liu (2013), among others. Despite the simplicity of this approach, it does not take into account the effect of the regularization term on classification function estimation. For binary problems, Zhang *et al.* (2013) showed that the regularization terms can lead to severe bias in class probability estimation. In particular, for high dimensional problems, we often need a heavy penalty in the corresponding optimization to ensure a good generalization ability of the resulting classifier. However, heavy penalties can introduce significant shrinkage effect on the fitted classification functions. Consequently, the estimated probabilities would be biased. To alleviate this difficulty, for binary margin-based classifiers, Zhang *et al.* (2013) proposed a two-stage refit method. The idea of Zhang *et al.* (2013) is that despite the bias in probability estimation, margin-based classifiers can often find near optimal classification boundaries. Consequently, one can project all observations in the training data set onto the norm vector of the separating hyperplane, and perform a second step classifier training using the projected coordinates as the only predictor. Because there is only one predictor in the refit step, the regularization term is no longer needed, thus it does not suffer from the shrinkage effect. Through numerical examples, Zhang *et al.* (2013) showed that the refit method works very well for binary problems, in the sense that the second step classifier has similar prediction accuracy compared with the original one, and the corresponding probability estimation can be significantly improved.

Despite the progress of obtaining accurate probability estimation for binary margin-based classifiers, the effect of regularization terms on probability estimation for multicategory problems remains unclear. In this paper, we fill the gap by showing the shrinkage effect of the penalty on probability estimation for multicategory classification, and propose a new multicategory refit method to alleviate this difficulty. Our contribution is two fold. First, we provide some theoretical insights on why heavy penalty is often needed for high dimensional problems, and show that the fitted classification functions would have small norms. This leads to that the estimated probabilities would shrink towards $1/k$ for a k-class problem. Our second contribution is that we propose a multicategory refit method for simultaneous margin-based classifiers. Note that for multicategory problems, because there are more than one separating boundaries, how to perform appropriate projection becomes more involved. We propose to use the $k-1$ classification functions introduced in Zhang and Liu (2014) as the coordinates for the second step classifier training, and show that our new method can significantly improve class probability estimation in empirical studies.

The rest of this article is organized as follows. In Section 2, we give a brief review of both binary and multicategory margin-based classifiers, and introduce how to estimate class conditional

probability. In Section 3, we first give some theoretical insights on why heavy penalty is often needed for high dimensional low sample size problems, then discuss the corresponding consequence on class probability estimation. Then, we propose our multicategory refit method to alleviate such difficulties. Numerical examples including simulation and real data analysis in Section 4 demonstrates the usefulness of our new method. Some concluding remarks are given in Section 5. All technical proofs are collected in the Supplementary Materials.

2 Review of Margin-based Classifiers

For a classification problem, suppose the training data set $\{(\boldsymbol{x}_i, y_i);\ i = 1, \ldots, n\}$ is obtained from an unknown underlying distribution $P(\boldsymbol{X}, Y)$, where $\boldsymbol{X} \in \mathbb{R}^p$ is a vector of predictors, and $Y \in \{+1, -1\}$ is the label for a binary problem. The typical learning goal is to obtain a classifier from the training data set, such that the future prediction error rate is minimized. In other words, we aim to find a prediction rule $\hat{y}(\cdot)$ that minimizes the expected 0-1 loss $E[I\{\hat{y}(\boldsymbol{X}) \neq Y\}]$, where $I(\cdot)$ is the indicator function. Binary margin-based classifiers use a single function $f(\boldsymbol{x})$ for classification, and the corresponding prediction rule is $\hat{y}(\boldsymbol{x}) = \text{sign}\{f(\boldsymbol{x})\}$. Consequently, minimizing the expected 0-1 loss is equivalent to minimizing $E[I\{Yf(\boldsymbol{X}) < 0\}]$. However, because the indicator function is non-convex and discontinuous, direct minimization of the empirical 0-1 loss can be difficult. To overcome this challenge, a common approach is to use a surrogate convex loss function ℓ in place of the indicator function. In particular, a typical binary margin-based classifier solves the optimization

$$\min_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} \ell\{y_i f(\boldsymbol{x}_i)\} + \lambda J(f),$$

where $\mathcal{F}$ is the functional space of f, $J(f)$ is the regularization term on f to prevent overfitting, and λ is a tuning parameter to balance the loss and regularization terms. In practice, an appropriate choice of λ can be critical for the performance of the resulting classifier. Note that for an observation $(\boldsymbol{x}, y)$, $yf(\boldsymbol{x})$, known as the functional margin, is positive if and only if the prediction is correct. Therefore, many binary margin-based classifiers use a non-increasing loss function ℓ to encourage large functional margins yf.

When the number of classes $k > 2$, it is common to have $Y \in \{1, \ldots, k\}$. In this case, one cannot use a single function f for classification. In the literature, many simultaneous margin-based classifiers use k functions $\boldsymbol{f} = (f_1, \ldots, f_k)$ for learning, and the corresponding prediction rule is

$\hat{y} = \text{argmax}_{j=1,\ldots,k} f_j$. A typical simultaneous classifier solves the following optimization

$$\min_{\boldsymbol{f} \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} \phi\{y_i, \boldsymbol{f}(\boldsymbol{x}_i)\} + \lambda J(\boldsymbol{f}),$$
$$\text{subject to } \sum_{j=1}^{k} f_j = 0, \tag{1}$$

where ϕ is a multicategory margin-based loss function. Here, to reduce the parameter space and to ensure theoretical properties such as Fisher consistency, the sum-to-zero constraint on $\boldsymbol{f}$, namely $\sum_{j=1}^{k} f_j(\boldsymbol{x}) = 0$ for all $\boldsymbol{x}$, is often imposed. Different loss functions correspond to different simultaneous margin-based classifiers. For example, the Multicategory SVM (MSVM) proposed by Lee *et al.* (2004) uses $\phi(y, \boldsymbol{f}) = \sum_{j \neq y} [1 + f_j]_+$, where $[u]_+ = \max(0, u)$, the MSVM by Liu and Shen (2006) uses $\phi(y, \boldsymbol{f}) = \sum_{j \neq y} [1 - \min_j (f_y - f_j)]_+$, and the multicategory large margin unified machine by Zhang and Liu (2013) uses $\phi(y, \boldsymbol{f}) = \gamma V(f_y) + (1 - \gamma) \sum_{j \neq y} V(-f_j)$, where $\gamma \in [0, 1]$ is a convex combination parameter, and V is the binary large margin unified machine loss function (see the appendix for more details). Note that these methods all encourage a large value of $f_{y_i}(\boldsymbol{x}_i)$, which can be regarded as the functional margin in multicategory problems.

Recently, Zhang and Liu (2014) pointed out that for a k-class problem, using k classification functions and the sum-to-zero constraint can be inefficient, and proposed an angle-based classification framework to overcome this challenge. The details of angle-based classifiers are as follows. Consider a centered simplex with k vertices $\boldsymbol{W} = \{\boldsymbol{W}_1, \ldots, \boldsymbol{W}_k\}$ in $\mathbb{R}^{k-1}$

$$\boldsymbol{W}_j = \begin{cases} (k-1)^{-1/2} \mathbf{1} & j = 1, \\ -(1 + k^{1/2})/\{(k-1)^{3/2}\} \mathbf{1} + \{k/(k-1)\}^{1/2} e_{j-1} & 2 \leq j \leq k, \end{cases}$$

where $\mathbf{1}$ is a vector of 1, and e_j is a vector with jth element 1 and 0 elsewhere. One can verify that $\boldsymbol{W}$ defines a symmetric simplex in $\mathbb{R}^{k-1}$. Without loss of generality, assume class j is assigned to $\boldsymbol{W}_j$. Zhang and Liu (2014) proposed to use $\boldsymbol{f}(\boldsymbol{x}) = \{f_1(\boldsymbol{x}), \ldots, f_{k-1}(\boldsymbol{x})\} : \mathbb{R}^p \to \mathbb{R}^{k-1}$ as the classification function vector. Each $\boldsymbol{f} \in \mathbb{R}^{k-1}$ defines k angles with respect to $\boldsymbol{W}_1, \ldots, \boldsymbol{W}_k$, namely, $\angle(\boldsymbol{W}_j, \boldsymbol{f})$; $j = 1, \ldots, k$. The prediction rule is $\hat{y}(\boldsymbol{x}) = \text{argmin}_{j \in \{1,\ldots,k\}} \angle(\boldsymbol{W}_j, \boldsymbol{f})$. Note that the smaller $\angle(\boldsymbol{W}_j, \boldsymbol{f})$, the larger the corresponding inner product $\langle \boldsymbol{W}_j, \boldsymbol{f} \rangle$. Therefore, the goal of the angle-based classifier is equivalent to maximizing $\langle \boldsymbol{W}_y, \boldsymbol{f}(\boldsymbol{x}) \rangle$, which can be regarded as functional margins in the angle-based classification. In particular, an angle-based classifier solves

the optimization problem

$$\min_{\boldsymbol{f}\in\mathcal{F}} \frac{1}{n}\sum_{i=1}^{n} \ell\{\langle \boldsymbol{W}_{y_i}, \boldsymbol{f}(\boldsymbol{x}_i)\rangle\} + \lambda J(\boldsymbol{f}), \tag{2}$$

where ℓ is a binary margin-based loss function. The angle-based classifiers are free of the sum-to-zero constraint, therefore can enjoy a more efficient computational speed. Zhang and Liu (2014) showed that angle-based classifiers are highly competitive for high dimensional classification problems.

Besides class prediction accuracy, estimation of class probabilities can be very important. For example, in clinical trials, physicians may have different treatments for patients based on their estimated probabilities of a certain disease. In the margin-based classification literature, to estimate the probabilities, a common approach is to use the fitted classification functions. In particular, we denote by $P_j(\boldsymbol{x}) = \text{pr}(Y = j \mid \boldsymbol{X} = \boldsymbol{x})$ the class conditional probability for class j. Define the conditional expected loss

$$S(\boldsymbol{x}) = E[\phi\{Y, \boldsymbol{f}(\boldsymbol{X})\} \mid \boldsymbol{X} = \boldsymbol{x}], \tag{3}$$

where ϕ is a margin-based loss function. Note that if $\boldsymbol{f}$ is one dimensional and one chooses ϕ appropriately, (3) can include binary classifiers as special cases. Next, define the theoretical minimizer of (3) to be

$$\boldsymbol{f}^*(\boldsymbol{x}) = \operatorname*{arginf}_{\boldsymbol{f}} S(\boldsymbol{x}) = \operatorname*{arginf}_{\boldsymbol{f}} E[\phi\{Y, \boldsymbol{f}(\boldsymbol{X})\} \mid \boldsymbol{X} = \boldsymbol{x}]. \tag{4}$$

For a classification problem with $Y \in \{1, \ldots, k\}$, under mild conditions (e.g., the loss function ϕ is differentiable), one can find the relationship between $P_j(\boldsymbol{x})$ and $\boldsymbol{f}^*(\boldsymbol{x})$, namely, $P_j(\boldsymbol{x}) = g_j\{\boldsymbol{f}^*(\boldsymbol{x})\}$ for all j. When the fitted classification function $\hat{\boldsymbol{f}}$ is obtained for a new observation $\boldsymbol{x}$, one can use $\hat{P}_j(\boldsymbol{x}) = g_j\{\hat{\boldsymbol{f}}(\boldsymbol{x})\}$ to estimate the probability of class j for $\boldsymbol{x}$. For example, in binary margin-based classification, under mild conditions, one can estimate $P_{+1}(\boldsymbol{x})$ by

$$\ell'(-\hat{f})/\{\ell'(-\hat{f}) + \ell'(\hat{f})\}, \tag{5}$$

where ℓ is the corresponding binary margin-based loss function (Zou *et al.*, 2008; Zhang *et al.*, 2013). For the angle-based classification, Zhang and Liu (2014) showed that the relationship between P_j

and $\boldsymbol{f}^*$ is such that

$$P_j = \frac{\ell'(\langle \boldsymbol{W}_j, \boldsymbol{f}^* \rangle)^{-1}}{\sum_{i=1}^{k} \ell'(\langle \boldsymbol{W}_i, \boldsymbol{f}^* \rangle)^{-1}}, \tag{6}$$

for all $j \in \{1, \ldots, k\}$ and differentiable loss function ℓ. Consequently, one can estimate P_j by $\hat{P}_j = \{\ell'(\langle \boldsymbol{W}_j, \hat{\boldsymbol{f}} \rangle)^{-1}\} / \{\sum_{i=1}^{k} \ell'(\langle \boldsymbol{W}_i, \hat{\boldsymbol{f}} \rangle)^{-1}\}$. Note that Wang *et al.* (2008) and Wu *et al.* (2010) proposed methods for class conditional probability estimation using SVM classifiers.

3 Methodology

In this section, we first explore the effect of the regularization term on class probability estimation in Section 3.1, then propose a new refit method for multicategory classification in Section 3.2.

3.1 Shrinkage Effect of Probability Estimation in Multicategory Classification

For binary margin-based classifiers, Zhang *et al.* (2013) observed that the L_2 regularization term can have a shrinkage effect on the estimated classification function, which can further lead to biased probability estimation. Despite the progress, many challenges remain. For example, whether L_1 type and other regularization terms such as the smoothly clipped absolute deviation (SCAD, Fan and Li, 2001) penalty have a similar shrinkage effect was not explored, and the consequence of regularization terms for multicategory classification is unclear. In this section, we extend the discussion to angle-based classifiers for various regularization methods, and demonstrate that the shrinkage effect is common for high dimensional problems.

To begin with, we first provide some theoretical insights on why heavy penalty is often needed for high dimensional low sample size classification problems. Recall that the primary goal of classification is to obtain a classifier with high prediction accuracy for future testing data sets. Traditional methods such as ordinary logistic regression minimize the empirical loss function on the training data set, and it is known that this approach can result in a classifier that fits the training data too well with suboptimal future prediction. For high dimensional problems, the overfitting issue is more severe. To overcome this difficulty, it is common to use regularization terms in the optimization to guarantee the generalization ability of the resulting classifier. In the next theorem, we focus on the angle-based classifier (2) with linear learning and L_1 and L_2 penalties, and show that in order to have a good classifier for problems with large p and small n, we often need λ to

be large. Note that $f_j(\boldsymbol{x}) = \boldsymbol{\beta}_j^T \boldsymbol{x} + \beta_{0,j}$ for all $j \in \{1, \ldots, k-1\}$ in linear learning. In this paper, we assume the domain of $\boldsymbol{X}$ is bounded, and the loss function ℓ is continuously differentiable. One can verify that this assumption is mild and can be satisfied by many practical problems.

Theorem 1. *Define $I\{y, \boldsymbol{f}(\boldsymbol{x})\}$ to be the indicator function of misclassification for the observation $(\boldsymbol{x}, y)$, using $\boldsymbol{f}(\cdot)$ and the angle-based prediction rule. We have that, with probability at least $1-\delta$,*

$$E[I\{Y, \hat{\boldsymbol{f}}(\boldsymbol{X})\}] \leq \frac{1}{n}\sum_{i=1}^{n} I\{y_i, \hat{\boldsymbol{f}}(\boldsymbol{x}_i)\} + 3\left\{\frac{\log(2/\delta)}{n}\right\}^{1/2} + Z,$$

where

$$Z = \begin{cases} \frac{C_1}{\lambda}\sqrt{\frac{\log(p)}{n}} & \text{for linear learning with } L_1 \text{ penalty,} \\ \frac{C_2}{\lambda}\sqrt{\frac{p}{n^{1/2}}} & \text{for linear learning with } L_2 \text{ penalty,} \end{cases}$$

and C_1, C_2 are constants independent of λ, n and p.

Theorem 1 shows that for a trained classification function vector $\hat{\boldsymbol{f}}$ from (2), the corresponding prediction error rate on a large testing data set can be bounded by the error rate on the training set, plus a penalty term that depends on λ, n and p. For a problem with large p and small n, one can verify that we need a large λ in order to guarantee that the future error rate is small. This helps to shed some light on why heavy penalty is often needed for high dimensional problems.

Remark 1. *For brevity, in Theorem 1 we focus on linear learning in angle-based classifiers using L_1 and L_2 regularization. One can verify that similar results hold for many other settings, for example learning in a Reproducing Kernel Hilbert Space (RKHS, Schölkopf and Smola, 2002; Shawe-Taylor and Cristianini, 2004), using other penalties such as the SCAD (Fan and Li, 2001) and the elastic-net (Zou and Hastie, 2005) methods, and classification with k functions and the sum-to-zero constraint (1). See the proof of Theorem 1 in the appendix for more details.*

Next, we show that the regularization term with large λ can significantly shrink the norm of $\hat{\boldsymbol{f}}$, which further leads to the bias in probability estimation using (6) for angle-based classifiers. As in Theorem 1, we focus on L_1 and L_2 penalties in linear learning in the following proposition.

Proposition 1. *The solution $\hat{\boldsymbol{f}}$ to (2) satisfies that $\|\hat{\boldsymbol{f}}\|_1 \leq C/\lambda$ for the L_1 and L_2 regularized learning, where C is a universal constant independent of λ, n and p.*

Because every element of $\boldsymbol{W}_j$ is bounded by 1, we have that $\langle \boldsymbol{W}_j, \hat{\boldsymbol{f}}(\boldsymbol{x})\rangle$ is bounded by C/λ for all j and $\boldsymbol{x}$. For high dimensional problems, Theorem 1 shows that the classifier tends to

use a large λ for the regularization term. As a result, the functional margins in the angle-based classification framework, $\langle \boldsymbol{W}_j, \hat{\boldsymbol{f}}(\boldsymbol{x}) \rangle$, tend to shrink towards zero. This, however, does not affect the label prediction performance of the classifier, as a correct prediction only requires that the corresponding angle-based functional margin $\langle \boldsymbol{W}_y, \hat{\boldsymbol{f}}(\boldsymbol{x}) \rangle$ is the largest. Therefore, with a large λ, the future prediction error rate $EI\{Y, \hat{\boldsymbol{f}}(\boldsymbol{X})\}$ can be close to the optimal (assuming that the empirical error rate on the training set is small), although the estimated $\hat{\boldsymbol{f}}$ has a small norm. On the other hand, the estimated probabilities tend to shrink towards $1/k$. In particular, recall from (6) that we use

$$\hat{P}_j(\boldsymbol{x}) = \frac{\ell'\{\langle \boldsymbol{W}_j, \hat{\boldsymbol{f}}(\boldsymbol{x}) \rangle\}^{-1}}{\sum_{i=1}^{k} \ell'\{\langle \boldsymbol{W}_i, \hat{\boldsymbol{f}}(\boldsymbol{x}) \rangle\}^{-1}} \tag{7}$$

to estimate $P_j(\boldsymbol{x})$. By Proposition 1, we have that $\langle \boldsymbol{W}_j, \hat{\boldsymbol{f}}(\boldsymbol{x}) \rangle \to 0$ for large λ. Consequently, $\ell'\{\langle \boldsymbol{W}_j, \hat{\boldsymbol{f}}(\boldsymbol{x}) \rangle\} \to \ell'(0)$. Therefore, for a high dimensional problem where the selected λ is large, $\hat{P}_j(\boldsymbol{x}) \to 1/k$ for all j and k.

Remark 2. *Using similar proofs as that of Proposition 1, one can verify that the shrinkage effect also exists for many other regularization methods, such as the SCAD and elastic-net penalties, and for learning in nonlinear functional spaces such as RKHS. See the proof and discussion of Proposition 1 in the appendix for more details.*

To alleviate the bias in probability estimation, in the next section, we propose a new refit method for angle-based classifiers. We show that our method can improve probability estimation accuracy, at the cost of an additional optimization problem that can be solved very efficiently.

3.2 A New Refit Method for Angle-based Classifiers

With appropriate regularization terms and tuning parameters in high dimensional problems, Theorem 1 and Proposition 1 show that margin-based classifiers can deliver accurate class prediction, however, the estimated functional margins could shrink towards 0. This can result in suboptimal probability estimation. To overcome this difficulty, in this section, we propose a new multicategory refit method that can significantly alleviate the bias in probability estimation.

To begin with, recall that it is common to perform probability estimation using (5) in binary problems. The following proposition shows that all the information about class conditional probability for binary margin-based classification is contained in the estimated $\hat{f}(\boldsymbol{x})$.

Proposition 2. *Suppose for two observations* $\boldsymbol{x}$ *and* $\boldsymbol{z}$*, we have that* $\hat{P}_1(\boldsymbol{x}) = \hat{P}_1(\boldsymbol{z})$ *using (5). Then, it must be that* $\hat{f}(\boldsymbol{x}) = \hat{f}(\boldsymbol{z})$*.*

Notice that the mapping $\boldsymbol{x} \to \hat{f}(\boldsymbol{x})$ is equivalent to finding the separating hyperplane in the feature space, and projecting all training observations onto its corresponding norm vector (see Figure 1 for an illustration). From Theorem 1 and Proposition 1, one can conclude that the sign of $\hat{f}(\boldsymbol{x})$ can be accurate, yet the scale would shrink towards 0 for high dimensional problems. To circumvent this challenge, Zhang *et al.* (2013) proposed to learn a second step classifier using $\{(\hat{f}(\boldsymbol{x}_i), y_i);\ i = 1, \ldots, n\}$ as the training data set, and employ probability estimation based on the result from the refit step. In other words, we employ the projected coordinates on the norm vector as the new variable. Zhang *et al.* (2013) proved that for binary classifiers where there is shrinkage effect, the refit step can help to "restore" the norm of estimated classification functions, which can lead to better probability estimation.

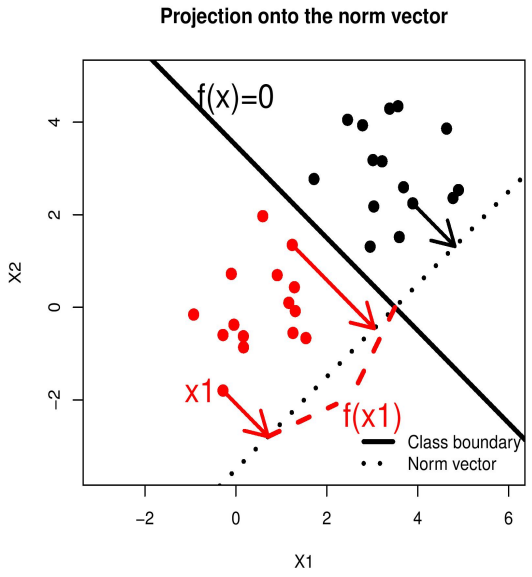

Figure 1: For binary margin-based classifiers, the classification function $f(\boldsymbol{x})$ maps $\boldsymbol{x}$ onto the norm vector of the separating hyperplane.

Despite the progress, how to perform the projection step in multicategory classification problems is challenging. For example, in the simultaneous margin-based classification with k functions and the sum-to-zero constraint, one can verify from similar arguments as in Theorem 1 and Proposition 1 that the label prediction based on $(\hat{f}_1, \ldots, \hat{f}_k)$ can be accurate, yet the probability estimation is biased because the norm of $(\hat{f}_1, \ldots, \hat{f}_k)$ shrinks. To alleviate this difficulty, a possible approach is

to use $(\hat{f}_1, \ldots, \hat{f}_k)$ as the second step input for classification without penalty terms. Similarly, in the angle-based classification framework, because class conditional probability estimation depends on $\langle \hat{\boldsymbol{f}}, \boldsymbol{W}_j \rangle$ for all j, one can use $\langle \hat{\boldsymbol{f}}, \boldsymbol{W}_j \rangle;\ j = 1, \ldots, k$ as the predictors for a naive refit approach. However, because of the commonly used sum-to-zero constraint for regular classifiers and the sum-to-zero property for angle-based methods, one of the k predictors in the second step refit approach mentioned above is unnecessary. It is unclear how to choose predictors that are not redundant, and simultaneously contain all information about class conditional probability. To overcome this challenge, for angle-based classifiers, we propose to use the $k-1$ dimensional $\{\hat{\boldsymbol{f}}(\boldsymbol{x}_i); i = 1, \ldots, n\} = \{(\hat{f}_1(\boldsymbol{x}_i), \ldots, \hat{f}_{k-1}(\boldsymbol{x}_i)); i = 1, \ldots, n\}$ as predictors for our second step refit method. The next theorem shows that the $k-1$ functions contain the same information about class conditional probabilities as $\langle \hat{\boldsymbol{f}}, \boldsymbol{W}_j \rangle;\ j = 1, \ldots, k$.

Theorem 2. *There is a one-to-one correspondence between $\boldsymbol{f} \in \mathbb{R}^{k-1}$ and $\langle \boldsymbol{f}, \boldsymbol{W}_j \rangle;\ j = 1, \ldots, k$, such that $\boldsymbol{f} = \boldsymbol{g}$ if and only if $\langle \boldsymbol{f}, \boldsymbol{W}_j \rangle = \langle \boldsymbol{g}, \boldsymbol{W}_j \rangle;\ j = 1, \ldots, k$. Moreover, if for two observations $\boldsymbol{x}$ and $\boldsymbol{z}$, we have that $(\hat{P}_1(\boldsymbol{x}), \ldots, \hat{P}_k(\boldsymbol{x})) = (\hat{P}_1(\boldsymbol{z}), \ldots, \hat{P}_k(\boldsymbol{z}))$ using (7), then $\hat{\boldsymbol{f}}(\boldsymbol{x}) = \hat{\boldsymbol{f}}(\boldsymbol{z})$.*

Because the norm of $\langle \hat{\boldsymbol{f}}, \boldsymbol{W}_j \rangle$ is an increasing function with respect to the norm of $\hat{\boldsymbol{f}}$, Theorem 2 suggests that if we can restore the norm of $\hat{\boldsymbol{f}}$, we can alleviate the shrinkage effect on class conditional probability estimation. We summarize our refit step for probability estimation in the angle-based classification framework as follows.

Algorithm for Multicategory Refit Method

1. Fit the classifier on the training data set to obtain $\hat{\boldsymbol{f}}$ with proper tuning on λ.

2. Create the training data set for the second step refit method. In particular, for the ith observation, we use $\hat{\boldsymbol{f}}(\boldsymbol{x}_i)$ as the corresponding predictor vector with $k-1$ elements, and y_i as the label.

3. Fit the classifier again on the training data set from Step 2, without any penalty terms. Denote by $\tilde{\boldsymbol{f}}$ the obtained classifier.

4. Estimate the class conditional probability $P_j(\boldsymbol{x})$ by $\tilde{P}_j(\boldsymbol{x}) = \frac{\ell'[\langle \boldsymbol{W}_j, \tilde{\boldsymbol{f}}\{\hat{\boldsymbol{f}}(\boldsymbol{x})\}\rangle]^{-1}}{\sum_{q=1}^{k} \ell'[\langle \boldsymbol{W}_q, \tilde{\boldsymbol{f}}\{\hat{\boldsymbol{f}}(\boldsymbol{x})\}\rangle]^{-1}}$.

Note that in Step 3 in the above algorithm, we train the refit classifier without regularization terms. This is because in most applications, the number of observations n is much larger than the number of classes k. Hence, there is little danger for overfitting, and no penalty is needed.

Consequently, there would be no shrinkage effect on the estimated $\tilde{\boldsymbol{f}}$, and we can obtain better probability estimation accuracy.

In the next section, we perform numerical analysis on simulated and real data sets to demonstrate the effectiveness of our refit approach for multicategory classification problems.

4 Numerical Examples

In this section, we conduct numerical experiments to demonstrate the effectiveness of our new refit method for multicategory classification. In particular, we perform simulation analysis in Section 4.1, and test on some real data sets in Section 4.2.

4.1 Simulation

We investigate the performance of our refit method using two simulated examples with high dimensional predictors in this section. In particular, we compare the class prediction accuracy, as well as probability estimation accuracy, for the traditional method and our new refit method. To measure the probability estimation accuracy, we report the mean absolute difference (MAD) between the true class conditional probability and the estimated value. The MAD is defined to be

$$\text{MAD} = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{k}|P_j(\boldsymbol{x}_i) - \hat{P}_j(\boldsymbol{x}_i)|.$$

For all examples and classifiers, we train the model on a training data set, and select the best tuning parameters from a candidate set $\{2^{-10}, 2^{-9}, \ldots, 2^9, 2^{10}\}$ that correspond to the smallest prediction error on a separate tuning data set. Then, we measure the class prediction and probability estimation accuracy on a large testing data set. For the choice of loss functions, we use the logistic regression deviance loss (logi) and the soft large-margin unified machine loss (soft, Liu *et al.*, 2011) for demonstration. We emphasize that our method can be generalized in a straightforward manner to other differentiable convex loss functions. We also train support vector machines using kernlab (KSVM) and report the corresponding classification accuracy as a comparison. Note that traditional SVM does not provide class conditional probability estimation. For each example, the prior probabilities for each class are the same.

Example 1. In this example, we let $k = 3$. The classification signal depends on 10 independent predictors $X_1, \ldots, X_{10}$. Given Y, the marginal distribution of each predictor is normal with variance

	Example 1		Example 2	
	Error	MAD	Error	MAD
Bayesian	5.51	-	35.0	-
Soft	5.58	15.4	36.9	13.9
Refit Soft	7.23	4.22	38.0	5.80
Logi	5.53	23.6	38.5	12.0
Refit Logi	7.13	3.94	43.0	8.81
KSVM	28.87	-	87.4	-

Table 1: The classification error rates (Error) and mean absolute differences (MAD) probability estimation for simulated examples 1 and 2.

$\sigma = 2$. For class 1, the mean vector of the normal distribution is $(3, 3, 3, 3, 0, 0, 0, 0, 0, 0)^T$. For class 2, the mean vector of the normal distribution is $(0, 0, 0, 3, 3, 3, 3, 0, 0, 0)^T$. For class 3, the mean vector of the normal distribution is $(0, 0, 0, 0, 0, 0, 3, 3, 3, 3)^T$. Next, we generalize noise predictors for high dimensional learning. In particular, we use 1490 variables, each following *i.i.d.* $N(0, 0.02)$. The training and tuning sets have size 300 each, and the testing set has 29400 observations.

Example 2. We generalize a 10-class problem for this example, where the classification signal depends on two predictors. The marginal distribution of X_1 and X_2 contains classification signal, and we generate another 498 variables that follow *i.i.d.* $N(0, 0.01)$. For each class, the marginal distribution of (X_1, X_2) is joint normal with constant variance matrix I_2, and the corresponding mean vectors are uniformly distributed on the circle $X_1^2 + X_2^2 = 9$.

The results for simulated Examples 1 and 2 are reported in Table 1. From the outcome, one can conclude that our refit method can significantly improve the accuracy of class conditional probabilities, compared to the original method proposed by Zhang and Liu (2014). Moreover, the corresponding classification accuracy of the original angle-based method is close to that of our refit ones. Hence, the refit approach is highly competitive. Note that the SVM performs poorly.

4.2 Real Data Analysis

In this section, we investigate the performance of our proposed refit method by applications on some real data sets. In particular, we apply the original angle-based classifier with probability estimation, as well as our refit method, on two benchmark data sets that can be found on the UCI machine learning repository. The first data set that we consider is Wine, which depicts the quality of 3 categories of wine with 13 variables. The second data set is Phishing websites, which contains 30 attributes and 2456 observations. More details for these data can be found at `https://archive.ics.uci.edu/ml/datasets/`.

	Wine		Phishing	
	Error	CRE	Error	CRE
Soft	52.9	5.85	47.0	2.03
Refit Soft	28.1	3.76	47.8	0.77
Logi	58.4	6.40	36.6	1.92
Refit Logi	31.6	5.79	42.3	1.29

Table 2: The classification error rates (Error) and cross entropy error (CRE) probability estimation for real data sets.

For both examples, we split the entire data into 6 parts. We use one of the parts as the testing set, and perform a four-fold cross validation on the remaining training data to select the best classifiers. The selection of tuning parameters is similar to that of the simulation examples. Since we do not know the underlying class conditional probability, we use the Cross Entropy Error (CRE, Zhang and Liu, 2014) to measure the goodness of fit of the estimated probabilities. In particular, for a testing set with n observations, the CRE is given as

$$\text{CRE} = -\frac{1}{n}\sum_{i=1}^{n}\hat{p}_{y_i}(\boldsymbol{x}_i)\log\{\hat{p}_{y_i}(\boldsymbol{x}_i)\}.$$

We report the average results for 100 replicates in Table 2. Notice that one can draw similar conclusions as in the simulation section, which demonstrates the effectiveness of our refit approach.

5 Discussion

In this paper, we investigate the effect of penalty terms for multicategory margin-based classifiers. In particular, we provide some theoretical insights on why heavy penalty is often needed for high dimensional low sample size classification problems. Then we show that such penalty can lead to severely biased class conditional probability. To overcome the difficulty, we propose a new multicategory refit approach for the angle-based classifiers. Numerical experiments show that the new method can improve class probability estimation with very low computation burden, hence can be very useful.

References

Bartlett, P. L., Jordan, M. I., and McAuliffe, J. D. (2006). Convexity, Classification, and Risk Bounds. *Journal of the American Statistical Association*, **101**, 138–156.

Boser, B. E., Guyon, I. M., and Vapnik, V. N. (1992). A Training Algorithm for Optimal Margin Classifiers. In D. Haussler, editor, *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, COLT '92, pages 144–152. Association for Computing Machinery, New York, NY, U.S.A.

Cai, T. and Liu, W. (2011). A Direct Estimation Approach to Sparse Linear Discriminant Analysis. *Journal of the American Statistical Association*, **106**(496), 1566–1577.

Cortes, C. and Vapnik, V. N. (1995). Support Vector Networks. *Machine Learning*, **20**, 273–297.

Fan, J. and Li, R. (2001). Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properties. *Journal of the American statistical Association*, **96**(456), 1348–1360.

Freund, Y. and Schapire, R. E. (1997). A Desicion-theoretic Generalization of On-line Learning and an Application to Boosting. *Journal of Computer and System Sciences*, **55**(1), 119–139.

Friedman, J. H., Hastie, T. J., and Tibshirani, R. J. (2000). Additive Logistic Regression: a Statistical View of Boosting. *Annals of Statistics*, **28**(2), 337–407.

Hastie, T. J., Tibshirani, R. J., and Friedman, J. H. (2009). *The Elements of Statistical Learning*. New York: Springer, 2nd edition.

Lee, Y., Lin, Y., and Wahba, G. (2004). Multicategory Support Vector Machines, Theory, and Application to the Classification of Microarray Data and Satellite Radiance Data. *Journal of the American Statistical Association*, **99**, 67–81.

Lin, X., Wahba, G., Xiang, D., Gao, F., Klein, R., and Klein, B. (2000). Smoothing Spline Anova Models for Large Data Sets with Bernoulli Observations and the Randomized GACV. *Annals of Statistics*, **28**(6), 1570–1600.

Liu, Y. (2007). Fisher Consistency of Multicategory Support Vector Machines. In *Eleventh International Conference on Artificial Intelligence and Statistics*, pages 289–296.

Liu, Y. and Shen, X. (2006). Multicategory ψ-learning. *Journal of the American Statistical Association*, **101**, 500–509.

Liu, Y. and Yuan, M. (2011). Reinforced Multicategory Support Vector Machines. *Journal of Computational and Graphical Statistics*, **20**(4), 901–919.

Liu, Y., Zhang, H. H., and Wu, Y. (2011). Soft or Hard Classification? Large Margin Unified Machines. *Journal of the American Statistical Association*, **106**, 166–177.

Park, S. Y., Liu, Y., Liu, D., and Scholl, P. (2010). Multicategory composite least squares classifiers. *Statistical Analysis and Data Mining*, **3**(4), 272–286.

Schölkopf, B. and Smola, A. J. (2002). *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond (Adaptive Computation and Machine Learning)*. The MIT Press.

Shawe-Taylor, J. S. and Cristianini, N. (2004). *Kernel Methods for Pattern Analysis*. Cambridge University Press, 1st edition.

Shen, X., Tseng, G. C., Zhang, X., and Wong, W. H. (2003). On ψ-learning. *Journal of the American Statistical Association*, **98**, 724–734.

Wang, J., Shen, X., and Liu, Y. (2008). Probability Estimation for Large Margin Classifiers. *Biometrika*, **95**(1), 149–167.

Wang, L. and Shen, X. (2007). On L_1-norm Multi-class Support Vector Machines: Methodology and Theory. *Journal of the American Statistical Association*, **102**, 595–602.

Wu, Y., Zhang, H. H., and Liu, Y. (2010). Robust Model-free Multiclass Probability Estimation. *Journal of the American Statistical Association*, **105**(489), 424–436.

Zhang, C. and Liu, Y. (2013). Multicategory Large-margin Unified Machines. *Journal of Machine Learning Research*, **14**, 1349–1386.

Zhang, C. and Liu, Y. (2014). Multicategory Angle-based Large-margin Classification. *Biometrika*, **101**(3), 625–640.

Zhang, C., Liu, Y., and Wu, Z. (2013). On the Effect and Remedies of Shrinkage on Classification Probability Estimation. *The American Statistician*, **67**(3), 134–142.

Zhu, J. and Hastie, T. J. (2005). Kernel Logistic Regression and the Import Vector Machine. *Journal of Computational and Graphical Statistics*, **14**(1), 185–205.

Zou, H. and Hastie, T. J. (2005). Regularization and Variable Selection via the Elastic Net. *Journal of the Royal Statistical Society, Series B,*, **67**(2), 301–320.

Zou, H., Zhu, J., and Hastie, T. J. (2008). New Multicategory Boosting Algorithms Based on Multicategory Fisher-Consistent Losses. *Annals of Applied Statistics*, **2**(4), 1290–1306.